

06/27/00
JC685 U.S. PTO

U.S. PTO
09/605421
06/27/00

PATENT
IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants: ALI et al.
Appl. No. : To Be Assigned Examiner: To Be Assigned
Filed : June 27, 2000 Art Unit: To Be Assigned
For : MICROPROCESSOR MEMORY SPACE ALLOCATION MANAGEMENT
Attorney Docket Number: 1.054US

REQUEST FOR FORMAL PATENT APPLICATION

Assistant Commissioner for Patents
Box PATENT APPLICATION
Washington, D.C. 20231

Sir:

Enclosed is a new patent application for filing today. The details regarding this application are as follows.

Title: MICROPROCESSOR MEMORY SPACE ALLOCATION
MANAGEMENT

Inventors: Saqib Ali
18533 Boysenberry Drive, Apt. 284
Gaithersburg, Maryland 20879
Citizenship: United States

Zoran Mladenovic
6116 Robinwood Road
Bethesda, Maryland 20817
Citizenship: United States

Bogdan Kosanovic
5904 Jarvis Lane
Bethesda, Maryland 20814
Citizenship: Yugoslavia

Attached is an application for patent including specification, claim, abstract of the disclosure, one sheet of informal drawing, and inventors' declaration.

09/605421

Request for Formal Patent Application
Attorney Docket No. 1.054US
June 27, 2000
Page 2

CALCULATION OF FEE

Basic Fee	\$ 690.00
Additional Fees:	
Number of independent claims is 1.	
Number of dependent claims is 0.	
Number of total claims is 1.	
<hr/>	
TOTAL FEES FOR THE ABOVE APPLICATION	\$ 690.00


Please charge \$690 for the filing fee to Deposit Account
Number 50-0964. A duplicate copy of this page is enclosed.

Please address all correspondence regarding this application
to the person named below:

Paul Grandinetti
c/o Telogy Networks, Inc.
20250 Century Boulevard
Germantown, Maryland 20874

Respectfully submitted,

27 June 2000
Date



Paul Grandinetti
Registration No. 30,754

Paul Grandinetti
c/o Telogy Networks, Inc.
20250 Century Boulevard
Germantown, Maryland 20874

(202) 429-4560

MICROPROCESSOR MEMORY SPACE ALLOCATION MANAGEMENT

* * * * *

CROSS-REFERENCE TO RELATED APPLICATIONS

5 Not applicable.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

Not applicable.

10 BACKGROUND OF THE INVENTION

002230-100000
This invention relates to memory management for microprocessors. In particular, the present invention relates to the management of scarce and sometimes overlapping memory space in microprocessor devices such as embedded processors and digital
15 signal processors.

Memory management for DSP engineers is a significant bottleneck in the development process. Currently it is a tedious process in which an engineer often draws a picture of the
20 available memory on a piece of paper and marks off the start addresses of where he/she would like to place the desired memory buffers. The list is made to determine the availability and conflict of the selection of memory locations. This process can be time consuming, tedious and error prone. Often times, when a
25 mistake is made, it can only be discovered through very intensive debugging which may take days. Furthermore, pieces of paper

often get lost, crumpled, etc. Sharing such pieces of paper with others (possibly in remote locations) may also be difficult. The need for a solution is urgent since these memory managing problems lie in the critical path of software development at this
5 time.

SUMMARY OF THE INVENTION

The present invention presents a tool that automates the task of mapping the memory buffers and heaps to physical space.
10 This tool is a Visual Memory Manager (VMM). Which has as its input a "recipe file" which is a memory buffer allocation table created by the user with a GUI linking tool such as Visual Linker sold by Texas Instruments as part of the Code Composer software set. This recipe file designates the locations, sizes and
15 overlays of all the buffers and heaps. The VMM checks the validity of the memory map specified in the recipe file. If the memory map is found to be invalid, the user is notified of the error. Otherwise, a memory table is created which serves as "hooks" for runtime buffer manipulation. This tool eliminates
20 the bottleneck of memory management and placement.

BRIEF DESCRIPTION OF THE DRAWINGS

For a better understanding of the nature of the present invention, reference is had to the following figures and detailed
25 description, wherein like elements are accorded like reference numerals, and wherein:

Figure 1 is a functional diagram illustrating the steps for creating a scan list of the allocated memory buffer space.

DETAILED DESCRIPTION OF PREFERRED EXEMPLARY EMBODIMENTS

5 It is common in a DSP build to have a large number of modules and buffers, sometimes several hundred in a single build. Each buffer will occupy memory space when the buffer is active. Data may be written or read from the memory space allocated to the active buffer. A simplified example is illustrated Table I
10 below. The first seven exemplary buffers, A - E and the first 101 memory address locations are illustrated and disc used, an actual program implementing the present invention will have a significantly larger number of buffers and will utilize significantly greater memory space. As illustrated in Table I,
15 buffer A for example may start at memory location 10 and run to memory location 30, buffer B may start at location 15 and run to location 62, as illustrated:

TABLE I

Buffer	A	B	C	D	E
Start Address	10	15	2	25	50
End Address	30	62	75	62	75

Once the address allocations for the buffers has been established by the programers, the present invention scans the memory allocations and addresses from the first memory address to the end of the available memory addresses and creates a scan
5 list, as illustrated in Figure 1.

The scanning produces an indicator each time a buffer is initially present and each time a buffer ceases to be present. For example, as the memory is scanned from at location 2, buffer
10 C becomes present and a notation is recorded, as the scan continues, the beginning of buffer A is noted at 10 and the beginning of buffer B is noted at 15. The end of buffer A is noted at 31, not 30, because buffer A is still present at memory location 30. The remainder of the buffer detection notations are
15 illustrated in Figure 1.

A link list, linking the start and end of each buffer to a specified memory address, is created. From the link list, a Ordered List indicating the existence of overlapping buffers is
20 created. The Ordered list for the example of Table I and Figure 1 is illustrated below in Table II

TABLE II

Memory Address	Buffers Present
2	C
10	A, C
15	A, B, C
25	A, B, C, D
31	A, B, C, D
50	B, C, D, E
63	B, C, D, E
76	C, E

Through establishment of the interval table, the conflicts can be readily and visibly ascertained by the programmer and managed. By tracking the beginning and ending of each buffer allocation, the present invention reduces the quantity of data needed for review. The programmer does not need to review all memory locations, just those location when a buffer starts and/or stops. By establishing the table II above, the programmer can visualize the conflicts. For example, a conflict exists between buffer B and buffer A beginning at memory location 15 and continuing until memory location 31. It is also easy to discern that a conflict between buffers D and E begins at memory location 50 and is resolved by memory location 63.

With respect to the implementation of the present invention in the application of the use of a Visual Linker GUI, the input to the VMM is a recipe file which the user will create using the Visual Linker GUI. As such, this GUI is used simply as means to
5 create this input. Although the GUI was not created with this purpose in mind, the present invention teaches how to use it in this way. The Visual Linker GUI is used to create an unplaced memory region and to place that region in physical memory.

10 For the VMM to correctly interpret the contents of the recipe file, the programmer declares conventions for specifying heaps, memory buffers, copy groups, etc with the Visual Linker GUI. The correct conventions specified herein are illustrative of a preferred exemplary embodiment. It will be apparent to one
15 skilled that conventions can be changed to suit a specific implementation without departing from the invention concept taught herein. For the purposes of the exemplary embodiment, the implementation must be followed to correctly specify the aforementioned components. The convention is defined to create
20 Visual Linker memory regions corresponding to different components of the exemplary memory management scheme.

Specifying Configuration Information:

To make judgements on the legality of buffer placement, heap placement and overlay specifications, the VMM will need some high level information regarding the entire build. This information
5 is referred to as "Configuration Information". The user specifies this information in the Visual Linker GUI by creating a special top-level configuration symbol named "CONFIGURATION" in the exemplary embodiment. Certain information is specified within the comment field of this symbol. The VMM will extract
10 this information and use it.

The configuration symbol provides the VMM with the following information:

15 The number of applications that exist.
 The application names.
 A mapping from application names to application IDs.
 The module names.
 A mapping from module names to module IDs.
20 Total number of module groups.
 For each module group, which application it is part of,
 what the member modules are and how many instances of
 this module group exist.

25 Grammar:

The grammar that the user must follow to specify the above information is shown below in Table III.. If this syntax is violated, the VMM will inform the user of an error.

30 The grammar in Table III is described in a BNF-like notation. The Symbol "::<=" represents equivalency. It can be read as "expands to". The Symbol "|" represents disjunction. It

can be read as "or". Symbols in Courier represent keywords. Symbols in <angle brackets> represent expressions that need further expansion according to the grammar. Symbols in *<italics-angle brackets>* represent expressions that need to be specified by the user. All white space in the configuration text is ignored. Anything between a "#" and the end of a line is interpreted as a comment and will also be ignored.

TABLE III

10	<Configuration>	::=	<Application ID Translation Table> <Module ID Translation Table> <List of Definition of Module-groups>
15	<Application ID Translation Table>	::=	Application Translation Table <List of Application Record> End Table
20	<List of Application Record>	::=	<Application Record> <List of Application Record><Application Record>
25	<Application Record>	::=	<Application Name> <Application ID>:
	<Application Name> length 30>	::=	<An alpha-numeric string of maximum
30	<Application ID>	::=	<An 8 bit unsigned integer in ANSI -C hexadecimal format>
	<Module ID Translation Table>	::=	Module Translation Table <List of Module Record> End Table
35	<List of Module Record>	::=	<Module Record> <List of Module Record><Module Record>
40	<Module Record>	::=	<Module Name> , <Module ID>;

```

<Module Name> ::= <An alpha-numeric string of maximum
length 30>

5 <Module ID> ::= <An 8 bit unsigned integer in ANSI
-C hexadecimal format>

<Module-Group Definition List> ::= <Module-Group Definition> |
<Module-Group Definition>
10 <Module-Group Definition List>
<Module-Group Definition> ::= Module-Group <Module-Group Name>
Application: <Application Name>;
<Instance specifier>
End Module-Group

15 <Module-Group Name> ::= <An alpha-numeric string of maximum
length 30>

<Instance Specifier> ::= <Instance count> instances of
20 <Module List>;

<Instance Count> ::= <An 8 bit unsigned integer in ANSI
-C decimal format>

25 <Module List> ::= <Module Name> | <Module
Name>,<Module List>

```

Example Configuration Symbol:

30 The text of an example configuration symbol is illustrated in Table IV. This text is a member of the language described by the grammar in TABLE III. This configuration is for exemplary purposes and may not accurately reflect the modules, applications and module groups in any existing build.

TABLE IV

```

# Now for the Application ID Translation table:
Application Translation Table
#      System      0x00;          # This is Assumed
40      Voice       0x01;
      Fax          0x02;
End Table

```

Now for the Module ID Translation table:

Module Translation Table

```
VCU      0x01;
VAU      0x02;
5  ECU      0x03;
PIU      0x04;
TDU      0x05;
VPU      0x06;
PVP      0x07;
10 TSU      0x08;
TTU      0x09;
CID      0x10;
ACU      0x11;
DPU      0x12;
15 FIU      0x13;
```

End Table

```
Module-Group Constant      # This module-group
    Application: System;    # is never
20    1 instances of ACU;    # overlaid.
End Module-Group
```

```
Module-Group PCMVoice      # Voice
    Application: Voice;
25    3 instances of VCU, VAU, ECU, PIU, TDU, VPU, PVP
        TSU, TTU, CID;
End Module-Group
```

```
Module-Group PacketVoice   # Voice
30    Application: Voice;
        2 instances of PVP, VPU;
End Module-Group
```

```
Module-Group FaxGroup      # Fax
35    Application: Fax;
        5 instances of PIU, DPU, FIU, ACU;
End Module-Group
```

Specifying Buffers:

40 Users of the VMM tool use the Visual Linker GUI to specify their buffer requirements. For each buffer they want to specify, they create a memory section in the Visual Linker GUI. This memory section is the GUI representation of the buffer being specified. This GUI object is referred to as the *buffer-memory-*

section. The VMM interprets each such buffer-memory-section as a buffer request.

The user sets the buffer-memory-section's name to

5 "CH<channel_ID>_<module group>_<module>_BUF<n>" where *n* is an in
sequence integer starting at 1 (refer to the definitions of
channel, module group and module if necessary). Its size is set
to the desired buffer size. The user sets the alignment
requirement for the buffer-memory-section and then places it in
10 physical memory. Additionally, the user may add up to four
"tags" to the buffer-memory-section's comment field on separate
lines:

"volatile" Presence of this tag indicates that this buffer is
volatile.

15 "start_addr_prog" - Presence of this tag indicates that the start
address is in program memory.

"ic_class_id = x" - Presence of this tag indicates that this
buffer's Intra-Channel Class ID is x.

"mem_class = y" - Presence of this tag indicates that this
20 buffer's Memory Class ID is y (legal values are DARAM, SARAM,
ERAM, IRAM (DARAM or SARAM) or HEAP).

The VMM has knowledge of certain relevant attributes of
buffers specified in this way. These attributes - and how they
25 are obtained from the GUI - are explained below:

Buffer Number:

An integer field that is extracted from the name of the buffer-memory-section that the user sets. This field is set to the *n* term from the name expression shown above.

5

Channel ID:

An integer field that is extracted from the name of the buffer-memory-section that the user sets. This field is set to the *channel_ID* term from the name expression shown above.

10

Module ID:

An integer field that is inferred from the name of the buffer-memory-section that the user sets. This field is set to the module ID taken from the Module ID Translation Table (see "Specifying Configuration Information") which corresponds to the *module* term from the name expression shown above.

15

Application ID:

An integer field that is inferred from the name of the buffer-memory-section that the user sets. This field is set to the application ID taken from the Application ID Translation Table (see "Specifying Configuration Information") which corresponds to the application that *module group* term from the name expression shown belongs to.

25

Size:

An integer field representing the worst-case size of the buffer in words. This field is set to the size that user set the buffer-memory-section to.

5

Memory Class:

An integer field with five possible values indicating whether this buffer should be placed in DARAM, SARAM, ERAM, IRAM (DARAM or SARAM) or HEAP. This field is extracted from the
10 comments of the buffer-memory-section. If no mem_class tag exists for this buffer-memory-section, the Memory Class will default to HEAP.

Intra-Channel Class ID:

15 A 30 character alphanumeric label. It indicates membership/non-membership in an Intra-Channel Class. This designation is used to determine which overlays are legal. (see Overlay Classes below). Those buffers (and only those buffers) having the same Intra-Channel Class ID are members of an Intra-
20 Channel Class. Members of an Intra-Channel Class must necessarily be a part of the same module, in the same module-group. This field is extracted from the comments of the buffer-memory-section. If no ic_class_id tag exists for this buffer-memory-section, the Intra-Channel Class ID will default to null (IE: not
25 a member of any Intra-Channel Class).

Start Address:

This long integer field will indicate the start address of this buffer during execution. This field is set to the start address of the buffer-memory-section that the user has placed in physical memory. It is important that the user already have correctly set the alignment requirements of this buffer-memory-section in order to get a valid placement. The VMM will not check for this.

10 Start Address Program Flag:

This Boolean flag indicates weather the Start Address field mentioned above is in program or data memory. This field is extracted from the comments of the buffer-memory-section. If a start_addr_prog tag exists for this buffer-memory-section then the Start Address Program Flag is set to true. Otherwise, it will default to false.

Volatile Flag:

This Boolean field will indicate weather or not this buffer is volatile. If a buffer is designated to be volatile, it is understood to be a scratch buffer and may be overlaid by any other buffer. A non-volatile designation means that it cannot be overlaid unless overridden by a specific overlay class designation (see below). This field is extracted from the comments of the buffer-memory-section. If a volatile tag exists

for this buffer-memory-section then the Volatile Flag is set to true. Otherwise, it will default to false.

Specifying Copy-Groups:

5 Copy-groups are collections of buffers that are copied/moved together. Every copy group has an execution and a store address. The *execution address* defines the location where the buffers is accessed from in real-time. The *store address* defines the location where the buffers is stored between real-time accesses.

10 Buffers within the same copy-group may not overlay each other. One should not use copy groups unless there are multiple groups that share the same execution address. The VMM will generate warnings for all copy groups that do not share their execution address with any other copy group.

15 Buffers belonging to different copy-groups may be overlaid regardless of their Volatility Flag field. Users of the VMM tool will use the Visual Linker GUI to specify their copy-group requirements. For each copy-group they want to specify, they

20 will create a memory section in the Visual Linker GUI. This memory section is the GUI representation of the copy-group being specified. We will call this GUI object the *cg-memory-section*. The user should set those buffer-memory-sections that correspond to buffers that are children of the copy-group in question to the

25 children of the *cg-memory-section*.

For example: If cg-memory-section A is an on-screen memory section in Visual Linker corresponding to copy-group B, then the user should set the buffer-memory-sections corresponding to all the members of A to be B's children.

5

The VMM will interpret each such cg-memory-section as a copy-group request. The user must set the cg-memory-section's name to "CG_<label>" where name can be any alphanumeric string of length 27. The user must set the alignment requirement for the cg-memory-section. Additionally, the user may add up to four tags to the cg-memory-section's comment field on separate lines:

- "one_way_copy" - Presence of this tag indicates only 1-way copying is required.
- 15 "exec_addr_prog" - Presence of this tag indicates that the execution start address is in program memory.
- "store_addr = x" - Presence of this tag indicates that this copy-group's store address is x.
- "store_addr_prog" - Presence of this tag indicates that the store
20 start address is in program memory.

Below is a list of the copy group parameters stored by the VMM with an explanation of how they are obtained from the GUI.

25

Copy Group Name:

A 30 character alpha-numeric string field that is extracted from the name of the cg-memory-section that the user sets. This field is set to the *label* term from the name expression shown
5 above.

Application ID:

An integer field which is set to the Application ID field of the first buffer-memory-section that is a child of this cg-
10 memory-section. All other buffer-memory-sections in this cg-memory-section must have the same Application ID. If not, the Visual Memory Manager tool will alert the user of an error.

Memory Class:

15 An integer field with five possible values indicating whether this buffer should be placed in DARAM, SARAM, ERAM, IRAM (DARAM or SARAM) or HEAP. This field is set to the Memory Class field of the first buffer-memory-section that is a child of this cg-memory-section. All other buffer-memory-sections in this cg-
20 memory-section must have the same Memory Class. If not, the Visual Memory Manager tool will alert the user of an error.

Size:

This integer indicates the size in words of the entire copy-
25 group. This field is set to the size of the cg-memory-section in the Visual Linker GUI. The Visual Linker will already have

computed the size of this cg-memory-section to be the sum of the size of all its children (and it will have adjusted appropriately for any alignment requirements of the children).

5 2-Way Copy Flag:

This field will indicate weather or not this buffer must be saved back outside of an internal memory region after it has been used. If this Boolean-value field is set, 2-way copying is implied. Otherwise 1-way copying is implied. An example of where
10 1-way copying may be useful is while using constants such as the codec table. This field is extracted from the comments of the cg-memory-section. If a one_way_copy tag exists for this cg-memory-section then the 2-Way Copy Flag field is set to false. Otherwise, it will default to true.

15

Execution Address:

This long integer is the address that the buffer is stored in when/if it is copied into internal memory. This field is set to the start address of the cg-memory-section in the Visual
20 Linker GUI. It is important that the user already have correctly set the alignment requirements of this cg-memory-section in order to get a valid placement. The VMM will not check for this.

25

Execution Address Program Flag:

This Boolean flag indicates weather the Execution Address field mentioned above is in program or data memory. This field is extracted from the comments of the cg-memory-section. If an
5 exec_addr_prog tag exists for this cg-memory-section then the Execution Address Program Flag field is set to true. Otherwise, it will default to false.

Store Address:

10 This long integer is the address that the buffer is stored in when/if it is copied out of internal memory. This field is extracted from the comments of the buffer-memory-section. If no store_addr tag exists for this cg-memory-section, the Store Address will default to Execution Address.

Store Address Program Flag:

15 This Boolean flag indicates weather the Store Address field mentioned above is in program or data memory. This field is extracted from the comments of the cg-memory-section. If a
20 store_addr_prog tag exists for this cg-memory-section then the Store Address Program Flag field is set to true. Otherwise, it will default to false.

Overlay Scheme:

The user may place certain buffers/heaps in the same location using the Visual Linker GUI. This implies an overlay. The VMM will check the validity of overlays specified by the user against the following criteria. If any of the criteria are satisfied, then the VMM will allow the overlay. Otherwise, it will inform the user of an error. All information that the VMM requires to check the validity of overlays is entered by the user when the buffer-memory-regions and configuration symbols are created (in the Visual Linker GUI). The overlay rules are:

Application Overlay Criterion:

Any buffers that have the same Channel ID but different Application ID may overlay each other. An example of this type of overlay may occur in a scenario when two different applications exist: Voice and Fax. Assuming that both of these applications occur on the same channel, the buffers used by FIU (application: Fax) may overlay the buffers used by ECU (application: Voice). A caveat to this rule exists: buffers belonging to the System application (which has predefined application ID 0) may not be overlaid by any other buffers.

Intra-Channel Overlay Criterion:

Any buffers that have the same Intra-Channel Class ID may overlay each other. An example of this type of overlay may occur in a scenario when switching between two different codecs on the

same channel during a call. Buffers of the two different codecs may overlay each other. Members of an Intra-Channel Class must be in the same module-group (and thus by definition also in the same application). If not, the VMM will inform the user of the error.

5

Scratch Overlay Criterion:

Any buffer which has its Volatile Flag field set can be freely overlaid by any other buffer which also has its Volatile Flag set. This type of overlay occurs in the scratch memory region.

10

Copy Group Criterion:

Any buffers that are members of different copy groups may overlay each other if the associated copy-groups have the same Application ID. The copy groups create the overlay at the execution address by construction. That is, the copy groups are intended to be used for generating such overlays that save internal memory space. For example, a typical copy group would have its execution address in internal RAM and the store address in external RAM (data or program). Another copy group within the same application would share the execution address if the program flow allows it. That is, the two groups are not accessed at the same time.

15

20

This type of overlay is similar to the Intra-Channel overlay except that in this case the execution and store addresses differ.

5 Specifying Heaps:

Certain buffers will not have statically pre-set base addresses. Instead their base addresses is allocated at run time from a designated heap. We call these special buffers *heap buffers*. Every module-group/channel combination must have a
10 unique heap associated with it. Every heap buffer requested is allocated from the heap belonging to that module-group/channel combination.

An example may help clarify this point: Suppose there are 5
15 channels on a specific build; a fax module-group (application: fax) which has 3 instances and a voice module-group (application: voice) which has 5 instances. Each module-group/channel combination will have its own heap. So a heap buffer requested by the ECU module on channel 4 is allocated from the heap belonging
20 to the voice-group on channel 4. Other modules from this channel and this voice group (and only those modules) will receive their allocations from that heap. All other requests for heap buffers is met with allocations from other heap(s).

the heap's size is automatically computed by the Visual Linker GUI. The VMM will ignore all information regarding the buffers, however, this method is recommended to minimize errors and keep the concepts of heap allocation clear to the user. If this method
5 is not adhered to, incorrectly sized heaps may result.

The present invention has been described in terms of implementation through the existing GUI Visual Linker, however, the present invention also teaches an GUI specifically for the
10 task of memory management. This GUI allows easier specification, duplication, manipulation and examination of memory entities (buffers, heaps, copy groups) than the current Visual Linker GUI allows. This GUI may have some abilities for automating the actual placement of the buffers and/or some facility for
15 informing users of the most MIPS-effective buffer placement scheme.

Because many varying and different embodiments may be made within the scope of the inventive concept herein taught, and
20 because many modifications may be made in the embodiments herein detailed in accordance with the descriptive requirements of the law, it is to be understood that the details herein are to be interpreted as illustrative and not in a limiting sense.

WE CLAIM:

1. A method for identification of memory address allocation conflicts of a plurality of finite buffers within a defined memory space, comprising the steps of:

5 entering a list of said buffers and corresponding memory address allocations;

scanning said memory allocations from a first memory address to a second memory address within said memory space;

10 creating a link list of primary memory addresses correlating to the start and end of each of said buffers;

creating an ordered list of said primary memory addresses and corresponding buffers which include said addresses from said primary address list.

15

* * * * *

ABSTRACT OF THE DISCLOSURE

A method for identification of memory assignment conflicts in the assignment of memory location addresses to a set of buffers. Programs run in embedded processors using buffers in a fixed storage space need to be mapped to addresses which do not overlap or create conflicts. The process of assigning start and end addresses for buffers can be tedious and error prone if performed without automation. The present invention presents a tool that automates the task of mapping the memory buffers and heaps to physical space. The tool utilizes a memory buffer allocation table created by the programmer. The table designates the locations, sizes and overlays of all the buffers and heaps. The tool checks the validity of the memory map specified. If it is found to be invalid, the user is notified of the error. Otherwise, a memory table is created which will serve as "hooks" for runtime buffer manipulation.

The diagram shows a vertical bar representing a memory table. The top is labeled 'address 0' and the bottom is labeled 'Address 101'. The bar is divided into sections labeled A, B, C, D, and E. Arrows indicate the mapping of addresses to data points. For example, address 2 points to C, address 10 points to A, address 15 points to B, address 25 points to D, address 31 points to A, address 50 points to E, address 63 points to B, D, and address 76 points to C, E.

Address	Data Point(s)
2	C
10	A
15	B
25	D
31	A
50	E
63	B, D
76	C, E

FIGURE 1

FIGURE 1

Paul Grandinetti
c/o Telogy Networks, Inc.
20250 Century Boulevard
Germantown, Maryland 20874 Telephone: (202) 429-4560

[illegible]

INVENTORS' DECLARATION
MICROPROCESSOR MEMORY SPACE
ALLOCATION MANAGEMENT
Page 2

CO-INVENTOR: Sagib Ali Citizenship: United States

CO-INVENTOR: Zoran Mladenovic Citizenship: United States

CO-INVENTOR: Bogdan Kosanovic Citizenship: Yugoslavia

Co-Inventor's signature: Tommy Boyer Date: 6/22/00
Residence and Post Office Address: 5904 Jarvis Lane, Bethesda, Maryland 20814